

WEST Search History

DATE: Friday, May 19, 2006

<u>Hide?</u>	<u>Set Name</u>	<u>Query</u>	<u>Hit Count</u>
<i>DB=PGPB,USPT,USOC; PLUR=YES; OP=ADJ</i>			
<input type="checkbox"/>	L16	L15 and (monitor\$ same performance)	46
<input type="checkbox"/>	L15	L14 and instruction cache	209
<input type="checkbox"/>	L14	L12 and (counting or counter or count or counted)	411
<i>DB=EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=ADJ</i>			
<input type="checkbox"/>	L13	L12	3
<i>DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=ADJ</i>			
<input type="checkbox"/>	L12	L11 and performance	559
<input type="checkbox"/>	L11	L7 and (instruction near2 range)	800
<input type="checkbox"/>	L10	L9 and counting	30
<input type="checkbox"/>	L9	L8 and monitor\$	109
<input type="checkbox"/>	L8	L7 and (contiguous range)	162
<input type="checkbox"/>	L7	identif\$ same instruction	88567
<input type="checkbox"/>	L6	L5 and (contiguous range)	1
<input type="checkbox"/>	L5	L4 and (instruction near2 range)	27
<input type="checkbox"/>	L4	L3 and monitor\$ and count\$	662
<input type="checkbox"/>	L3	717/124-135.ccls.	2498
<input type="checkbox"/>	L2	L1 and (contiguous range).clm.	11
<input type="checkbox"/>	L1	(identif\$ same instruction).clm.	25669

END OF SEARCH HISTORY

[Sign in](#)
[Web](#) [Images](#) [Groups](#) [News](#) [Froogle](#) [Maps](#) [more »](#)
 [Search](#) [Advanced Search](#) [Preferences](#)
WebResults 1 - 10 of about 729 for [monitoring performance instruction "contiguous range" counting](#). (0.36 seconds)**Optimizing the Instruction Cache Performance of the Operating System**

The **performance monitor** cannot capture instruction accesses that hit in the ... The table also shows that the static **count** of the **instructions** in these ...

doi.ieeecomputersociety.org/10.1109/12.737683 - [Similar pages](#)

[PDF] Optimizing Instruction Cache Performance for Operating System ...

File Format: PDF/Adobe Acrobat

The **performance monitor** cannot capture **instruction** ac- cases that hit in the first level ... **count** of the **instructions** in these loops is as low as around ...

doi.ieeecomputersociety.org/10.1109/HPCA.1995.386527 - [Similar pages](#)

[[More results from doi.ieeecomputersociety.org](#)]

[PDF] Optimizing the Instruction Cache Performance of the Operating System

File Format: PDF/Adobe Acrobat - [View as HTML](#)

The **performance monitor** cannot capture **instruction** ac- ... shows that the dynamic **count** of the **instructions** in these. loops is only 29-39 percent of all ...

iacoma.cs.uiuc.edu/iacoma-papers/osplacelong.pdf - [Similar pages](#)

[PDF] 494 494..505

File Format: PDF/Adobe Acrobat - [View as HTML](#)

However, the **performance monitor** cannot. capture **instruction** accesses that hit in the primary ... **contiguous range** of memory locations called the SelfConf- ...

iacoma.cs.uiuc.edu/iacoma-papers/comprehen.pdf - [Similar pages](#)

[PDF] HPCView: A Tool for Top-down Analysis of Node Performance

File Format: PDF/Adobe Acrobat - [View as HTML](#)

Performance monitoring can be made more efficient by making use of special hardware features beyond simple. event **counting**. For instance, the experimental ...

www.cs.rice.edu/~mgabi/papers/hpcview-lacsi01.pdf - [Similar pages](#)

[PS] HPCView: A Tool for Top-down Analysis of Node Performance John ...

File Format: Adobe PostScript - [View as HTML](#)

5 Related Work **Performance monitoring** can be made more efficient by making use of special hardware features beyond simpleevent **counting**. ...

www.cs.rice.edu/~mgabi/papers/hpcview-lacsi01.ps - [Similar pages](#)

[[More results from www.cs.rice.edu](#)]

[PDF] Nios II Software User Guide

File Format: PDF/Adobe Acrobat - [View as HTML](#)

clear, or refresh the **count** values. 3.1. Enabling **Performance Monitor** Hardware ... The next PM example shows **counting** load. **instruction** misses in the Nios ...

www.fs2.com/nios_download/Nios2-SW-UserGuide.pdf - [Similar pages](#)

HP OpenVMS Technical Journal V6

On an Alpha system, BUGCHECK turns off **performance monitoring**. ... A fragment is a **contiguous range** of pages with common attributes (for example, ...

h71000.www7.hp.com/openvms/journal/v6/fatal_bugchecks.html - 93k -

[Cached](#) - [Similar pages](#)

[PDF] Abstract 1 Introduction

File Format: PDF/Adobe Acrobat - [View as HTML](#)

degradation over the base system, indicating that the. penalty on the web server **performance** is small. We further analyze this experiment by **monitoring** the ...

www.stanford.edu/class/cs240/readings/navarro.pdf - [Similar pages](#)

[PDF] [Optimizing Network Virtualization in Xen](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

the number of Xen **instruction** TLB misses. (SP-GP vs. SP). The overall impact of using global mappings on the. transmit **performance**, however, is not very ...

infoscience.epfl.ch/getfile.py?recid=85598&mode=best - [Similar pages](#)

Try your search again on [Google Book Search](#)

Gooooooooogle ►
Result Page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [Next](#)

Free! Speed up the web. [Download the Google Web Accelerator](#).

[Search within results](#) | [Language Tools](#) | [Search Tips](#) | [Dissatisfied? Help us improve](#)

[Google Home](#) - [Advertising Programs](#) - [Business Solutions](#) - [About Google](#)

©2006 Google

 **PORTAL**
USPTO

Subscribe (Full Service) Register (Limited Service, Free) Login
 Search: The ACM Digital Library The Guide
 "performance counter" +contiguous

THE ACM DIGITAL LIBRARY  Feedback Report a problem Satisfaction survey

Terms used performance counter contiguous

Found 2,134 of 176,279

Sort results by

 relevance Save results to a Binder

Try an Advanced Search

Display results

 expanded form Search TipsTry this search in The ACM Guide Open results in a new window

Results 21 - 40 of 200

Result page: previous

[1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale **21 Positional adaptation of processors: application to energy reduction** 

Michael C. Huang, Jose Renau, Josep Torrellas

May 2003 **ACM SIGARCH Computer Architecture News , Proceedings of the 30th annual international symposium on Computer architecture ISCA '03**, Volume 31 Issue 2

Publisher: ACM Press

Full text available:  pdf(225.57 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

Although adaptive processors can exploit application variability to improve performance or save energy, effectively managing their adaptivity is challenging. To address this problem, we introduce a new approach to adaptivity: the *Positional* approach. In this approach, both the *testing* of configurations and the *application* of the chosen configurations are associated with particular code sections. This is in contrast to the currently-used *Temporal* approach to adaptation ...

22 Procedure placement using temporal-ordering information 

Nikolas Gloy, Michael D. Smith

September 1999 **ACM Transactions on Programming Languages and Systems (TOPLAS)**, Volume 21 Issue 5

Publisher: ACM Press

Full text available:  pdf(604.56 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Instruction cache performance is important to instruction fetch efficiency and overall processor performance. The layout of an executable has a substantial effect on the cache miss rate and the instruction working set size during execution. This means that the performance of an executable can be improved by applying a code-placement algorithm that minimizes instruction cache conflicts and improves spatial locality. We describe an algorithm for procedure placement, one type of code placement ...

Keywords: code placement, conflict misses, temporal profiling, working-set optimization

23 Memory systems: Cluster miss prediction with prefetch on miss for embedded CPU instruction caches 

Ken Batcher, Robert Walker

September 2004 **Proceedings of the 2004 international conference on Compilers, architecture, and synthesis for embedded systems**

Publisher: ACM Press

Full text available:  pdf(343.66 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Soft CPU cores are often used in embedded systems, yet they limit opportunities to improve cache performance to hardware assistance outside the CPU core. Instruction

prefetching is commonly used, but the popular Prefetch On Miss (POM) technique is less helpful when the instruction flow does not follow a sequential execution order, which is often the case in real-time networking applications. Cluster Miss Prediction (CMP) can help in those worst case situations when cache misses do not follow a s ...

Keywords: WCET, cache design, cache prefetch, embedded systems, hiding memory latency, networking

24 Data page layouts for relational databases on deep memory hierarchies 

Anastassia Ailamaki, David J. DeWitt, Mark D. Hill

November 2002 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 11 Issue 3

Publisher: Springer-Verlag New York, Inc.

Full text available:  pdf(593.86 KB) Additional Information: [full citation](#), [abstract](#), [index terms](#)

Relational database systems have traditionally optimized for I/O performance and organized records sequentially on disk pages using the N-ary Storage Model (NSM) (a.k.a., slotted pages). Recent research, however, indicates that cache utilization and performance is becoming increasingly important on modern platforms. In this paper, we first demonstrate that in-page data placement is the key to high cache performance and that NSM exhibits low cache utilization on modern platforms. Next, we ...

Keywords: Cache-conscious database systems, Disk page layout, Relational data placement

**25 Compiling parallel languages: An evaluation of global address space languages: co-
array fortran and unified parallel C** 

Cristian Coarfa, Yuri Dotsenko, John Mellor-Crummey, François Cantonnet, Tarek El-Ghazawi, Ashrujit Mohanti, Yiyi Yao, Daniel Chavarría-Miranda

June 2005 **Proceedings of the tenth ACM SIGPLAN symposium on Principles and practice of parallel programming**

Publisher: ACM Press

Full text available:  pdf(246.41 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Co-array Fortran (CAF) and Unified Parallel C (UPC) are two emerging languages for single-program, multiple-data global address space programming. These languages boost programmer productivity by providing shared variables for inter-process communication instead of message passing. However, the performance of these emerging languages still has room for improvement. In this paper, we study the performance of variants of the NAS MG, CG, SP, and BT benchmarks on several modern architectures to iden ...

Keywords: CAF, UPC, co-array fortran, compilers, global address space languages, parallel languages, performance, scalability, unified parallel C

26 Cross-Platform Performance Prediction of Parallel Applications Using Partial Execution 

Leo T. Yang, Xiaosong Ma, Frank Mueller

November 2005 **Proceedings of the 2005 ACM/IEEE conference on Supercomputing SC '05**

Publisher: IEEE Computer Society

Full text available:  pdf(462.49 KB) Additional Information: [full citation](#), [abstract](#)
 Publisher Site

Performance prediction across platforms is increasingly important as developers can choose from a wide range of execution platforms. The main challenge remains to perform accurate predictions at a low-cost across different architectures. In this paper, we derive an affordable method approaching cross-platform performance translation based on

relative performance between two platforms. We argue that relative performance can be observed without running a parallel application in full. We show that ...

27 Scaling irregular parallel codes with minimal programming effort

 Dimitrios S. Nikolopoulos, Constantine D. Polychronopoulos, Eduard Ayguadé
November 2001 **Proceedings of the 2001 ACM/IEEE conference on Supercomputing (CDROM)**

Publisher: ACM Press

Full text available:  pdf(153.05 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The long foreseen goal of parallel programming models is to scale parallel code without significant programming effort. Irregular parallel applications are a particularly challenging application domain for parallel programming models, since they require domain specific data distribution and load balancing algorithms. From a performance perspective, shared-memory models still fall short of scaling as well as message-passing models in irregular applications, although they require less coding effort ...

28 Research sessions: implementation techniques: Implementing database operations using SIMD instructions

 Jingren Zhou, Kenneth A. Ross
June 2002 **Proceedings of the 2002 ACM SIGMOD international conference on Management of data SIGMOD '02**

Publisher: ACM Press

Full text available:  pdf(1.39 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Modern CPUs have instructions that allow basic operations to be performed on several data elements in parallel. These instructions are called SIMD instructions, since they apply a single instruction to multiple data elements. SIMD technology was initially built into commodity processors in order to accelerate the performance of multimedia applications. SIMD instructions provide new opportunities for database engine design and implementation. We study various kinds of operations in a database con ...

29 Session II: memory systems: An empirical performance analysis of commodity memories in commodity servers

 Darren J. Kerbyson, Mike Lang, Gene Patino, Hossein Amidi
June 2004 **Proceedings of the 2004 workshop on Memory system performance**

Publisher: ACM Press

Full text available:  pdf(207.86 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This work details a performance study of six different types of commodity memories in two commodity server nodes. A number of micro-benchmarks are used that measure low-level performance characteristics, as well as two applications representative of the ASC workload. The memories vary both in terms of performance, including latency and bandwidths, and in terms of their physical properties and manufacturer. The two server nodes analyzed were an Itanium-II Madison based system, and a Xeon based sy ...

Keywords: memory modules, memory system performance, performance analysis, performance measurement

30 Resource partitioning in general purpose operating systems: experimental results in Windows NT

 D. G. Waddington, D. Hutchison
October 1999 **ACM SIGOPS Operating Systems Review**, Volume 33 Issue 4

Publisher: ACM Press

Full text available:  pdf(1.56 MB) Additional Information: [full citation](#), [abstract](#), [index terms](#)

The principal role of the operating system is that of resource management. Its task is to

present a set of appropriate services to the applications and users it supports. Traditionally, general-purpose operating systems, including Windows NT, federate resource sharing in a fair manner, with the predominant goal of efficient resource utilisation. As a result the chosen scheduling algorithms are not suited to applications that have stringent Quality-of-Service (QoS) and resource management require ...

31 Hardware: Impact of modern memory subsystems on cache optimizations for stencil computations 

 Shoaib Kamil, Parry Husbands, Leonid Oliker, John Shalf, Katherine Yelick
June 2005 **Proceedings of the 2005 workshop on Memory system performance MSP '05**

Publisher: ACM Press

Full text available:  pdf(618.02 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

In this work we investigate the impact of evolving memory system features, such as large on-chip caches, automatic prefetch, and the growing distance to main memory on 3D stencil computations. These calculations form the basis for a wide range of scientific applications from simple Jacobi iterations to complex multigrid and block structured adaptive PDE solvers. First we develop a simple benchmark to evaluate the effectiveness of prefetching in cache-based memory systems. Next we present a small ...

Keywords: cache blocking, performance modeling, prefetch, stencil

32 Profile-based dynamic voltage and frequency scaling for a multiple clock domain microprocessor 

 Grigoris Magklis, Michael L. Scott, Greg Semeraro, David H. Albonesi, Steven Dropsho
May 2003 **ACM SIGARCH Computer Architecture News , Proceedings of the 30th annual international symposium on Computer architecture ISCA '03**, volume 31 Issue 2

Publisher: ACM Press

Full text available:  pdf(541.83 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

A Multiple Clock Domain (MCD) processor addresses the challenges of clock distribution and power dissipation by dividing a chip into several (coarse-grained) clock domains, allowing frequency and voltage to be reduced in domains that are not currently on the application's critical path. Given a reconfiguration mechanism capable of choosing appropriate times and values for voltage/frequency scaling, an MCD processor has the potential to achieve significant energy savings with low performance degr ...

33 A hardware-driven profiling scheme for identifying program hot spots to support runtime optimization 

 Matthew C. Merten, Andrew R. Trick, Christopher N. George, John C. Gyllenhaal, Wen-mei W. Hwu
May 1999 **ACM SIGARCH Computer Architecture News , Proceedings of the 26th annual international symposium on Computer architecture ISCA '99**, volume 27 Issue 2

Publisher: IEEE Computer Society, ACM Press

Full text available:  pdf(349.69 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index](#)

 Publisher Site

terms

This paper presents a novel hardware-based approach for identifying, profiling, and monitoring hot spots in order to support runtime optimization of general purpose programs. The proposed approach consists of a set of tightly coupled hardware tables and control logic modules that are placed in the retirement stage of a processor pipeline removed from the critical path. The features of the proposed design include rapid detection of program hot spots after changes in execution behavior, runtime-tu ...

34 Reconfigurable computing: architectures and applications: Using reconfigurability to achieve real-time profiling for hardware/software codesign 

-  Lesley Shannon, Paul Chow
 February 2004 **Proceedings of the 2004 ACM/SIGDA 12th international symposium on Field programmable gate arrays**
 Publisher: ACM Press
 Full text available: [pdf\(228.02 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Embedded systems combine a processor with dedicated logic to meet design specifications at a reasonable cost. The attempt to amalgamate two distinct design environments introduces many problems, one being how to partition a single design for the two platforms to achieve the best performance with the least effort. Since the latest FPGA technology allows the integration of soft or hard CPU cores with dedicated logic on a single chip, this presents new opportunities for addressing hardware/software ...

Keywords: FPGA, embedded processor, hardware/software codesign, performance measurement, profiling, soft processor

35 A fast and accurate framework to analyze and optimize cache memory behavior 

-  Xavier Vera, Nerina Bermudo, Josep Llosa, Antonio González
 March 2004 **ACM Transactions on Programming Languages and Systems (TOPLAS)**, Volume 26 Issue 2
 Publisher: ACM Press

Full text available: [pdf\(270.06 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#), [review](#)

The gap between processor and main memory performance increases every year. In order to overcome this problem, cache memories are widely used. However, they are only effective when programs exhibit sufficient data locality. Compile-time program transformations can significantly improve the performance of the cache. To apply most of these transformations, the compiler requires a precise knowledge of the locality of the different sections of the code, both before and after being transformed. Cache ...

Keywords: Cache memories, optimization, sampling

36 Contributions: focus: new visualization techniques: Rivet: a flexible environment for computer systems visualization 

-  Robert Bosch, Chris Stolte, Diane Tang, John Gerth, Mendel Rosenblum, Pat Hanrahan
 February 2000 **ACM SIGGRAPH Computer Graphics**, Volume 34 Issue 1
 Publisher: ACM Press

Full text available: [pdf\(1.25 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

Rivet is a visualization system for the study of complex computer systems. Since computer systems analysis and visualization is an unpredictable and iterative process, a key design goal of Rivet is to support the rapid development of interactive visualizations capable of visualizing large data sets. In this paper, we present Rivet's architecture, focusing on its support for varied data sources, interactivity, composition and user-defined data transformations. We also describe the challenges of i ...

37 The Jalapeño dynamic optimizing compiler for Java 

-  Michael G. Burke, Jong-Deok Choi, Stephen Fink, David Grove, Michael Hind, Vivek Sarkar, Mauricio J. Serrano, V. C. Sreedhar, Harini Srinivasan, John Whaley
 June 1999 **Proceedings of the ACM 1999 conference on Java Grande**

Publisher: ACM Press

Full text available: [pdf\(1.34 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

38 Summary of the sigmetrics symposium on parallel and distributed processing 

Jeffrey K. Hillingsworth, Barton P. Miller

March 1999 **ACM SIGMETRICS Performance Evaluation Review**, Volume 26 Issue 4



Publisher: ACM Press

Full text available: [pdf\(1.17 MB\)](#)

Additional Information: [full citation](#), [index terms](#)

39 [Comparing the memory system performance of the HP V-class and SGI Origin 2000 multiprocessors using microbenchmarks and scientific applications](#)



Ravi Iyer, Nancy M. Amato, Lawrence Rauchwerger, Laxmi Bhuyan

May 1999 **Proceedings of the 13th international conference on Supercomputing**

Publisher: ACM Press

Full text available: [pdf\(1.13 MB\)](#)

Additional Information: [full citation](#), [references](#), [index terms](#)

40 [Predicting data cache misses in non-numeric applications through correlation profiling](#)



Todd C. Mowry, Chi-Keung Luk

December 1997 **Proceedings of the 30th annual ACM/IEEE international symposium on Microarchitecture**

Publisher: IEEE Computer Society

Full text available: [pdf\(876.36 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)



[Publisher Site](#)

To maximize the benefit and minimize the overhead of software-based latency tolerance techniques, we would like to apply them precisely to the set of dynamic references that suffer cache misses. Unfortunately, the information provided by the state-of-the-art cache miss profiling technique (summary profiling) is inadequate for references with intermediate miss ratios - it results in either failing to hide latency, or else inserting unnecessary overhead. To overcome this problem, we propose and ev ...

Keywords: profiling, cache miss prediction, correlation, non-numeric applications, latency tolerance.

Results 21 - 40 of 200

Result page: [previous](#) [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads: [Adobe Acrobat](#) [QuickTime](#) [Windows Media Player](#) [Real Player](#)